



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/671,889	09/29/2003	Fred Gehrunig Gustavson	YOR920030170US1	8009
48150 7590 07/16/2007 MCGINN INTELLECTUAL PROPERTY LAW GROUP, PLLC 8321 OLD COURTHOUSE ROAD SUITE 200 VIENNA, VA 22182-3817			EXAMINER VICARY, KEITH E	
			ART UNIT 2183	PAPER NUMBER
			MAIL DATE 07/16/2007	DELIVERY MODE PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary

Application No.

10/671,889

Applicant(s)

GUSTAVSON ET AL.

Examiner

Keith Vicary

Art Unit

2183

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 19 June 2007.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-20 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-20 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|---|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____ |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | 5) <input type="checkbox"/> Notice of Informal Patent Application |
| 3) <input checked="" type="checkbox"/> Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date <u>2/27/2007, 3/30/2007, 5/3/2007</u> . | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

1. Claims 1-20 are pending in this office action and presented for examination.

Claims 1-6 and 8-19 are amended and claim 20 is added by amendment filed 6/19/2007.

Double Patenting

2. Claims 1-20 of this application conflict with claims 1, 3-6, 8-12, and 14-19 of Application No. 10671937. 37 CFR 1.78(b) provides that when two or more applications filed by the same applicant contain conflicting claims, elimination of such claims from all but one application may be required in the absence of good and sufficient reason for their retention during pendency in more than one application. Applicant is required to either cancel the conflicting claims from all but one application or maintain a clear line of demarcation between the applications. See MPEP § 822.

3. The nonstatutory double patenting rejection is based on a judicially created doctrine grounded in public policy (a policy reflected in the statute) so as to prevent the unjustified or improper timewise extension of the "right to exclude" granted by a patent and to prevent possible harassment by multiple assignees. A nonstatutory obviousness-type double patenting rejection is appropriate where the conflicting claims are not identical, but at least one examined application claim is not patentably distinct from the reference claim(s) because the examined application claim is either anticipated by, or would have been obvious over, the reference claim(s). See, e.g., *In re Berg*, 140 F.3d 1428, 46 USPQ2d 1226 (Fed. Cir. 1998); *In re Goodman*, 11 F.3d 1046, 29 USPQ2d 2010 (Fed. Cir. 1993); *In re Longi*, 759 F.2d 887, 225 USPQ 645 (Fed. Cir. 1985); *In re Van Ornum*, 686 F.2d 937, 214 USPQ 761 (CCPA 1982); *In re Vogel*, 422 F.2d 438, 164 USPQ 619 (CCPA 1970); and *In re Thorington*, 418 F.2d 528, 163 USPQ 644 (CCPA 1969).

A timely filed terminal disclaimer in compliance with 37 CFR 1.321(c) or 1.321(d) may be used to overcome an actual or provisional rejection based on a nonstatutory double patenting ground provided the conflicting application or patent either is shown to

Art Unit: 2183

be commonly owned with this application, or claims an invention made as a result of activities undertaken within the scope of a joint research agreement.

Effective January 1, 1994, a registered attorney or agent of record may sign a terminal disclaimer. A terminal disclaimer signed by the assignee must fully comply with 37 CFR 3.73(b).

4. Claims 1-9 and 11-19 are provisionally rejected on the ground of nonstatutory obviousness-type double patenting as being unpatentable over claims 1, 3-6, 8-12, and 14-19 of copending Application No. 10671937. Although the conflicting claims are not identical, they are not patentably distinct from each other because claims 1-9 and 11-19 of the instant application are obvious variants of claims 1, 3-6, 8-12, and 14-19 of the '937 application.

This is a provisional obviousness-type double patenting rejection because the conflicting claims have not in fact been patented.

5. Claims 1-9 and 11-19 of the instant application contain every limitation of claims 1, 3-6, 8-12, and 14-19 of the '937 application; moreover, claims 1-9 and 11-19 of the instant application claim prefetching data into a cache providing data into an FPU, whereas claims 1, 3-6, 8-12, and 14-19 of the '937 application merely claim preloading data into a floating point register of an FPU.

It would have been readily recognized by one of ordinary skill in the art at the time of the invention that the benefits of using cache in the instant application are numerous and include greater system performance due to the decreased access time to access cache in comparison to main memory combined with the locality of reference that is typical in most computer programs.

Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention to implement cache into the instant application to gain greater system performance; it would have been readily recognized by one of ordinary skill in the art at the time of the invention that greater system performance is desirable in any processor. Furthermore, it would have been readily recognized by one of ordinary skill in the art at the time of the invention that this cache would fit into the '937 application by receiving data from the main memory and sending it to the floating point register, and that when preloading data into the floating point register in a system which uses a cache, that data would have to be prefetched into the cache in order to be preloaded into the register.

Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention to combine the widely-known teachings of cache with the invention of the '937 application in order to increase system performance.

a. Further note that claims 2, 11, and 13 in the instant application also claim that prefetching data is accomplished by utilizing time slots caused by a difference between a time to execute instructions in said subroutine execution process and a time to load said data, while claims 1, 11, and 12 of the '937 application does not explicitly disclose this.

It would have been readily recognized by one of ordinary skill in the art at the time of the invention that prefetching data in general cuts down the amount of time a processor is waiting for a memory miss to be serviced, and prefetching by utilizing time slots caused by a difference between a time to execute instructions

and a time to load said data allows for data to be prefetched ahead of time without delaying any other instructions that are being processed. Furthermore, it would have been readily recognized by one of ordinary skill in the art at the time of the invention that the benefits of prefetching are contingent upon other instructions not being delayed due to the prefetching; thus, it would have been readily recognized to one of ordinary skill in the art at the time of the invention that prefetching would be done by utilizing these time slots of inactivity.

Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention to combine the widely-known method of prefetching by utilizing time slots with the '937 application in order to cut down the amount of time a processor is waiting for a memory miss to be serviced, thus increasing overall system performance.

6. Aside from the obvious variants listed above, claim 1 of the '937 application contains every element of claim 1 of the instant application.
7. Aside from the obvious variants listed above, claim 1 of the '937 application contains every element of claim 2 of the instant application.
8. Aside from the obvious variants listed above, claim 3 of the '937 application contains every element of claim 3 of the instant application.
9. Aside from the obvious variants listed above, claim 4 of the '937 application contains every element of claim 4 of the instant application.

10. Aside from the obvious variants listed above, claim 5 of the '937 application contains every element of claim 5 of the instant application.
11. Aside from the obvious variants listed above, claim 6 of the '937 application contains every element of claim 6 of the instant application.
12. Aside from the obvious variants listed above, claim 8 of the '937 application contains every element of claim 7 of the instant application.
13. Aside from the obvious variants listed above, claim 9 of the '937 application contains every element of claim 8 of the instant application.
14. Aside from the obvious variants listed above, claim 10 of the '937 application contains every element of claim 9 of the instant application.
15. Aside from the obvious variants listed above, claim 6 of the '937 application contains every element of claim 11 of the instant application.
16. Aside from the obvious variants listed above, claim 12 of the '937 application contains every element of claim 12 of the instant application.
17. Aside from the obvious variants listed above, claim 12 of the '937 application contains every element of claim 13 of the instant application.
18. Aside from the obvious variants listed above, claim 14 of the '937 application contains every element of claim 14 of the instant application.
19. Aside from the obvious variants listed above, claim 15 of the '937 application contains every element of claim 15 of the instant application.
20. Aside from the obvious variants listed above, claim 16 of the '937 application contains every element of claim 16 of the instant application.

21. Aside from the obvious variants listed above, claim 17 of the '937 application contains every element of claim 17 of the instant application.

22. Aside from the obvious variants listed above, claim 18 of the '937 application contains every element of claim 18 of the instant application.

23. Aside from the obvious variants listed above, claim 19 of the '937 application contains every element of claim 19 of the instant application.

24. Amended claims besides claims 10 and 20 do not affect the overall double patenting rejection. It is noted that the claims of the '937 application which have been amended since the previous office action do not render the rejection obsolete.

b. Amended claims such as claim 1 include the addition of the limitation of improving efficiency. However, it would have been obvious to one of ordinary skill in the art at the time of the invention that prefetching data could improve efficiency.

c. Amended claims such as claim 1 include the addition of the limitation of a Level 3 Dense Linear Algebra Subroutine; however, as explained below, this limitation is encompassed by the '937's application disclosure of the level 3 BLAS.

d. Amended claims such as 1 and 6 replace the limitations of prefetching and touching with timely moving data by inserting moving instructions; it would have been obvious to one of ordinary skill in the art at the time of the invention that

prefetching with touch instructions is an example of inserting moving instructions to timely move data.

25. Claims 10 and 20 are provisionally rejected on the ground of nonstatutory obviousness-type double patenting as being unpatentable over claims 11 and 1 of copending Application No. 10671937 as explained above, and further in view of Gustavson et al. (Gustavson) (Superscalar GEMM-based Level 3 BLAS – The On-going Evolution of a Portable and High-Performance Library, Para'98, pages 207-215). This is a provisional obviousness-type double patenting rejection.

26. In addition to the limitations disclosed above in the rejection of claims 1 and 6, Claims 10 and 20 of the instant application additionally disclose a compiler *modified to incorporate linear algebra theory and techniques to automatically* generate instructions. The '937 application does not disclose this.

On the other hand, Gustavson does disclose of modifying a compiler to incorporate linear algebra theory and techniques to automatically generate instructions (modifying the compiler to incorporate linear algebra theory and techniques with the line "Unfortunately, todays compilers do not perform prefetching as well as one would desire, especially for complex memory hierarchies" in section 4.1, line 6-7. Given the word "unfortunately," it would have been readily recognized to one of ordinary skill in the art at the time of the invention that compilers performing prefetching as well as one would desire for complex memory hierarchies is being disclosed, such as prefetching such as the algorithmic prefetching disclosed in section 4.1, line 8. Note that the

algorithmic prefetching being cited incorporates linear algebra theory and techniques, as can be in the reference cited in the instant prior art and additionally cited in the additional references section below).

It would have been readily recognized to one of ordinary skill in the art at the time of the invention that a compiler which can insert touch instructions in optimal places and the like instead of making a human do it manually would lead to increased system performance at little to no cost to human capital.

Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention to combine the teachings of Gustavson with the invention of the '937 application in order to increase system performance at little to no cost to human capital.

Claim Rejections - 35 USC § 112

27. The following is a quotation of the first paragraph of 35 U.S.C. 112:

The specification shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same and shall set forth the best mode contemplated by the inventor of carrying out his invention.

28. The following is a quotation of the second paragraph of 35 U.S.C. 112:

The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

29. Claims 1-20 are rejected under 35 U.S.C. 112, first paragraph, as failing to comply with the written description requirement. The claim(s) contains subject matter which was not described in the specification in such a way as to reasonably convey to

one skilled in the relevant art that the inventor(s), at the time the application was filed, had possession of the claimed invention.

30. Claim 1-2, 6, 12-13, and 17 as amended recite the limitation “timely move” instead of “prefetch” or a variation thereof. Although prefetching data can certainly be thought of as timely moving data, the limitation “timely move” is broader than the limitation “prefetch”; in other words, there are scenarios or interpretations in which “timely moving data” is not synonymous with “prefetching data.” The specification on page 23, second paragraph, or page 4, first paragraph, discloses of prefetching data prior to the time that the data is actually required for the kernel calculations, for example, but “timely moving data” when read broadly could still allow for data to be moved into a cache after the time that the data is actually required for the kernel calculations, as long as the moving of data is not exceedingly long after.

e. Claims 2-5, 7-11, 13-16, and 18-20 are rejected for failing to alleviate the rejection of claims 1, 6, 12, and 17 above.

31. Claims 2, 6, 10-11, and 13 as amended recite the limitation “scheduling move type instructions into time slots” or a variation thereof. Similar to the rejection above, although the disclosed instructions to touch data can certainly be thought of as move type instructions, the limitation “move type instructions” is broader than data touch instructions; in other words, there are scenarios or interpretations in which “move type instructions” is not synonymous with data touch instructions. A review of the instant specification and the specifications incorporated by reference does not seem to yield a disclosed broad “move type instruction.”

- f. Claims 7-11 are rejected for failing to alleviate the rejection of claim 6 above.

32. Claim 2 as amended recites the limitation "scheduling move type instructions into time slots" in lines 1-2. However, a perusal of the instant specification and the specifications incorporated by reference does not seem to yield any explicit scheduling. The original claim discloses of merely having data touch instructions in open time slots; however, the amended claim discloses of *scheduling* these instructions into the time slots, which changes the scope of the claim and does not appear to be in any of the specifications.

33. Claims 2, 11, and 13 as amended recites the limitation "existing in a Level 3 Dense Linear Algebra Subroutine" or a variation thereof. The limitation does not appear to be present and connected to the other limitations in the instant or incorporated specifications and thus is considered new matter. If this Level 3 Dense Linear Algebra Subroutine is a subset of Level 3 BLAS, then the scope of the claim is being narrowed with no basis in the specification. If this Level 3 Dense Linear Algebra Subroutine is synonymous with Level 3 BLAS or something else in the specifications, then a reference or citation should be provided which validates this.

34. Claims 10 and 20 as amended recites the limitation "a compiler as modified to incorporate linear algebra theory and techniques." The limitations that describe a modified compiler do not appear to be present in the instant or incorporated specifications and thus is considered new matter.

35. Claims 10 and 20 as amended recites the limitation "a compiler as modified to incorporate linear algebra theory and techniques to automatically generate instructions." The limitation that the compiler *automatically* generates instructions does not appear to be present in the instant or incorporated specifications and thus is considered new matter. The limitation narrows the claim as, for example, a compiler may generate instructions for said inserting with human guidance and not completely automatically.

36. Claims 10 and 20 as amended recite the limitation "a compiler...to automatically generate instructions for said inserting said moving instructions" or a variation thereof. The limitation that there are generated instructions which themselves cause the insertion of additional moving instructions does not appear to be present in the instant or incorporated specifications and thus is considered new matter. See the 112 second rejection below regarding this limitation for more explanation:

g. Claim 11 is rejected for failing to alleviate the rejections of claim 10 above.

37. Claims 1-20 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

38. Claims 1, 4, 8, 12, 15, and 17-18 recite the limitation "thereby improving an efficiency" or a variation thereof. It is indefinite as to what is meant by this limitation. As one example, one definition of efficiency is "ability to accomplish a job with a minimum expenditure of time and effort." However, while it appears that the instant invention does minimize time spent in executing the linear algebra subroutine due to the

prefetching, the processor must consequently execute more instructions in order to perform the prefetching, which correlates to the processor expending more “effort” and perhaps power in executing the linear algebra subroutine. Thus, while time is minimizing, “effort” is increasing. Because there is no disclosed or commonly known way of determining at what point “efficiency” increases given both the decrease in time yet the increase in effort, the use of “efficiency” is indefinite. Examiner recommends amending to focus on, for example, the decrease in time needed to perform the subroutine execution.

- h. Claims 2-5, 13-16, and 18-20 are rejected for failing to alleviate the rejection of claims 1, 12, and 17 above.

39. Claim 6 recites the limitation “wherein said matrix data in said memory is timely moved by inserting moving instructions to be loaded into said cache” in lines 7-8. The limitation as written seems to imply that it is the moving instructions which are being loaded into said cache and not said matrix data, and should thus be rewritten to be more clear.

- i. Claims 7-11 are rejected for failing to alleviate the rejection of claim 6 above.

40. Claims 10 and 20 as amended recites the limitation “a compiler as modified to incorporate linear algebra theory and techniques.” This limitation is indefinite as the limitations “modified to incorporate linear algebra theory and techniques” does not particularly point out how *exactly* the compiler is modified.

41. Claims 10 and 20 as amended recite the limitation "a compiler...to automatically generate instructions for said inserting said moving instructions" or a variation thereof. However, it is indefinite as to what exactly the instructions for inserting instructions are. In contrast, the original claim disclosed generating instructions to touch data, as opposed to, for example, generating *instructions for* inserting instructions to touch data.

j. Claim 11 is rejected for failing to alleviate the rejections of claim 10 above.

Claim Rejections - 35 USC § 101

42. 35 U.S.C. 101 reads as follows:

Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

43. Claims 12-16 are rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter.

k. Functional descriptive material must be claimed in combination with an appropriate computer readable medium in order to be statutory.

44. Claim 12 is not limited to tangible embodiments. In view of Applicant's disclosure, specification page 20, lines 9-15, the signal-bearing medium is not limited to tangible embodiments, instead being defined as including both tangible embodiments (e.g., DASD storage, magnetic tape, electronic read-only memory, an optical storage device, paper "punch cards") and intangible embodiments (e.g., other suitable signal-bearing media including transmission media such as digital and analog and communication links and wireless). As such, the claim is not limited to statutory subject matter and is therefore non-statutory.

I. Claims 13-16 are rejected for failing to alleviate the rejection of claim 12 above.

45. To alleviate this rejection, the examiner recommends amending the specification to clearly label the aforementioned tangible embodiments as types of storage medium and the aforementioned intangible embodiments as types of transmission medium, and then subsequently amend the claim to disclose the storage medium instead of the signal-bearing medium.

46. This rejection is presently maintained for two reasons. The first reason is that the specification was not labeled to *clearly* label the signal-bearing media. As it is currently written, the use of the phrase "or other suitable signal-bearing media" implies that the signal-bearing media is a subset of the machine-readable data storage media.

Examiner recommends that, as explained above, the distinction be made more clear, such as by saying "the instructions may be stored on a variety of machine-readable data storage media, such as..., or as an alternative to machine-readable data storage media, on suitable signal bearing media including..." or something along those lines. Moreover, the use of the limitation "computer-readable storage medium" in the claims must have basis in the specification. Therefore, examiner recommends that that the aforementioned limitation be amended to read "machine-readable data storage medium" instead.

Claim Rejections - 35 USC § 102

47. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

48. Claims 1-20 are rejected under 35 U.S.C. 102(b) as being anticipated by Gustavson.

49. Consider claim 1, Gustavson discloses a method of executing a linear algebra subroutine, said method comprising: for an execution code (section 1, line 6, BLAS code) controlling an operation of a floating point unit (FPU) (section 3.1, line 4, discloses floating point registers, therefore it is inherent there are floating point units that are doing the multiplications as in section 1, line 2) performing a linear algebra subroutine execution (section 1, line 8, routine along with section 1, line 1, linear algebra), inserting instructions to timely data (section 4.1, lines 7-9, algorithmic prefetching) into a cache (section 4.1, line 4, cache) providing data into said FPU (section 4.1, line 1, data, and section 4.1, line 10, BLAS, which uses the FPUs), thereby improving an efficiency for said linear algebra subroutine execution (it is inherent that prefetching data during a subroutine execution may improve the execution time of a program as opposed to not prefetching data during a subroutine execution).

50. Consider claim 6, Gustavson discloses an apparatus, comprising: a memory to store matrix data to be used for processing in a linear algebra program (section 4, line

Art Unit: 2183

12, shared main memory and section 4.2, lines 7-9, elements of the matrix); a floating point unit (FPU) to perform said processing (section 3.1, line 4, discloses floating point registers, therefore it is inherent there are floating point units that are doing the multiplications as in section 1, line 2); a load/store unit (LSU) to load data to be processed by said FPU (section 3.1, lines 6-7, load and store operations, thus it is inherent there is a load/store unit), said LSU loading said data into a plurality of floating point registers (FRegs) (section 3.1, line 4, floating point registers); and a cache to store data from said memory and provide said data to said Fregs (section 4.1, line 4, cache), wherein said matrix data in said memory is timely moved by inserting moving instructions to be loaded into said cache prior to a need for said data to be in said FRegs for said processing, (section 4.1, line 12, touch instruction; lines 7-9, algorithmic prefetching).

51. Consider claim 12, Gustavson discloses a computer readable storage medium tangibly embodying a program of machine-readable instructions executable by a digital processing apparatus to perform a method of executing linear algebra subroutines, said method comprising: for an execution code (section 1, line 6, BLAS code) controlling an operation of a floating point unit (FPU) (section 3.1, line 4, discloses floating point registers, therefore it is inherent there are floating point units that are doing the multiplications as in section 1, line 2) performing a linear algebra subroutine execution (section 1, line 8, routine along with section 1, line 1, linear algebra), inserting instructions to timely move data (section 4.1, lines 7-9, algorithmic prefetching) into a

cache (section 4.1, line 4, cache) providing data into said FPU (section 4.1, line 1, data, and section 4.1, line 10, BLAS, which uses the FPUs), thereby improving an efficiency for said linear algebra subroutine execution (it is inherent that prefetching data during a subroutine execution may improve the execution time of a program as opposed to not prefetching data during a subroutine execution).

52. Consider claim 17, Gustavson discloses a method of providing a service involving at least one of solving and applying a scientific/engineering problem, said method comprising at least one of:

using a linear algebra software package that computes one or more matrix subroutines, wherein said linear algebra software package generates an execution code (section 1, line 6, BLAS code) controlling an operation of a floating point unit (FPU) (section 3.1, line 4, discloses floating point registers, therefore it is inherent there are floating point units that are doing the multiplications as in section 1, line 2) performing a linear algebra subroutine execution (section 1, line 8, routine along with section 1, line 1, linear algebra), unrolling such that instructions (section 3.1, line 1, loop unrolling) are inserted to timely move data (section 4.1, line 4, prefetching, lines 7-9, algorithmic prefetching) into a cache (section 4.1, line 4, cache) providing data for said FPU (section 4.1, line 1, data, and section 4.1, line 10, BLAS, which uses the FPUs), thereby improving an efficiency (it is inherent that prefetching data during a subroutine execution may improve the execution time of a program as opposed to not prefetching data during a subroutine execution) for said linear algebra subroutine execution (section 4.1, line

12, touch); providing a consultation for solving a scientific/engineering problem using said linear algebra software package (it is inherent that the BLAS will solve some type of scientific/engineering problem for someone who may or may not be the operator of the BLAS program); transmitting a result of said linear algebra software package on at least one of a network, a signal-bearing medium containing machine-readable data representing said result, and a printed version representing said result; and receiving a result of said linear algebra software package on at least one of a network, a signal-bearing medium containing machine-readable data representing said result, and a printed version representing said result (it is inherent that the result of the problem will be conveyed to someone who may or may not be the operator of the BLAS program; furthermore, it is inherent that the result can only be shown either through a printout or through some type of electronic means, which encompasses voice through a phone or data through a network that is read via a monitor).

53. Consider claims 2 and 13, Gustavson discloses said timely moving is accomplished by inserting/scheduling move type instructions into time slots existing in a Level 3 Dense Linear Algebra Subroutine (Section 1, line 1, BLAS). As explained above, it is inherent to prefetching that data is loaded into the cache before the instruction that needs that data is executed, thus there must be a difference between the time of that instruction execution and the time of its data loading, otherwise it would not be prefetching. Furthermore, Gustavson discloses in page 212, lines 2-3 of section 4.1 that the prefetching instruction does not disturb ongoing computations and data

references, thus this prefetching must be done in "time slots" which are independent of other instruction fetching.

54. Consider claim 11, Gustavson discloses said moving instructions are inserted into time slots existing in a Level 3 Dense Linear Algebra Subroutine (Section 1, line 1, BLAS). As explained above, it is inherent to prefetching that data is loaded into the cache before the instruction that needs that data is executed, thus there must be a difference between the time of that instruction execution and the time of its data loading, otherwise it would not be prefetching. Furthermore, Gustavson discloses in page 12, lines 2-3 of section 4.1 that the prefetching instruction does not disturb ongoing computations and data references, thus this prefetching must be done in "time slots" which are independent of other instruction fetching.

55. Consider claims 3, 7, and 14, Gustavson discloses said matrix subroutine comprises a matrix multiplication operation (section 1, line 2, matrix multiply).

56. Consider claims 4, 8, 15, and 18, Gustavson discloses said matrix subroutine comprises a more efficient equivalent (it is inherent that prefetching data during a subroutine execution may improve the execution time of a program as opposed to not prefetching data during a subroutine execution; section 4.1, line 4, prefetching, lines 7-9, algorithmic prefetching) of a subroutine from a LAPACK (Linear Algebra PACKage) (section 1, line 1, discloses a BLAS, which is a part of LAPACK).

57. Consider claims 5, 9, 16, and 19, Gustavson discloses said LAPACK/linear algebra/processing subroutine invokes a BLAS Level 3 L1 cache kernel (Abstract, lines 1-6, level 3 BLAS kernel and level 1 cache).

58. Consider claims 10 and 20, Gustavson discloses a compiler to automatically generate instructions for said inserting said moving instructions (section 4.1, lines 2-4, compiler). Gustavson also discloses modifying the compiler to incorporate linear algebra theory and techniques with the line "Unfortunately, todays compilers do not perform prefetching as well as one would desire, especially for complex memory hierarchies" in section 4.1, line 6-7. Given the word "unfortunately," it would have been readily recognized to one of ordinary skill in the art at the time of the invention that compilers performing prefetching as well as one would desire for complex memory hierarchies is being disclosed, such as prefetching such as the algorithmic prefetching disclosed in section 4.1, line 8. Note that the algorithmic prefetching being cited incorporates linear algebra theory and techniques, as can be in the reference cited in the instant prior art and additionally cited in the additional references section below.

Response to Arguments

59. Applicant's arguments filed 6/19/2007 have been fully considered but they are not persuasive.

60. Applicant argues in general of the double patenting rejection. Examiner believes that the claim revisions in the present application are still obvious variants of the '937 application and is described in more detail above in the section on double patenting. Examiner maintains that pre-fetching is an obvious variant of preloading as well, as is described in this and the previous office action.

61. In response to applicant's request for reconsideration regarding the 101 rejection, examiner refers applicant to the section of the 101 rejection above.

62. Applicant argues on pages 11-12 that there are elements of the claimed invention that are not taught or suggested by this earlier publication by Gustavson. While there may be elements in the *overall disclosure* that are not taught or suggested by this earlier publication by Gustavson, there does not appear to be elements in the *claimed* invention that are not taught or suggested by this earlier publication by Gustavson. For example, the added limitation of "improving an efficiency" is met by the earlier publication as it would have been readily recognized that the algorithmic prefetching could be more "efficient" than typical or no prefetching. As another example, the *concept* of a modified compiler to incorporate linear algebra theory and techniques to automatically generate touch instructions *is* disclosed by the earlier publication by Gustavson. Although any *more specific* way of incorporating linear algebra theory and techniques to do this may not be disclosed by the earlier publication by Gustavson, it is not claimed in the instant set of claims. To provide an analogous example, the mere idea of modifying an airplane to account for newer discoveries in aerodynamics so that the airplane is more efficient and can travel further is not and

probably will never be novel; it is *what* the modification actually is which may be novel.

In a likewise manner, the claims will most likely not be novel unless specific details already located in the specification are brought into the claims.

63. In the context of the arguments at the bottom of page 11, the mere concept of scheduling pre-fetch instructions is not novel; however, a specific method of doing said scheduling may be novel. The concept of using knowledge about linear algebra in order to make compilers produce code which executes faster than normal is not novel; however, a specific method or delineation of changes of how a compiler produces this code which executes faster than normal may be novel. It is once again noted that all of these potentially novel points must be claimed and not merely in the specification.

Conclusion

64. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

m. Gustavson et al. (Improving performance of linear algebra algorithms for dense matrices, using algorithmic prefetch) further details the practice of algorithmic prefetching. Page 266, section prefetching, also in essence discloses the concept of a compiler inserting prefetch instructions for complex programs. Although the disclosure acknowledges that the compiler cannot or with great difficulty judiciously insert prefetch instructions, it nevertheless still discloses the concept. Although the instant application may teach a specific way in which

compilers can judiciously insert prefetch instructions, this specific way must be claimed

65. Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.

66. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Keith Vicary whose telephone number is (571) 270-1314. The examiner can normally be reached on Monday - Friday, 8:00 a.m. - 5:00 p.m., EST.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Eddie Chan can be reached on 571-272-4162. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Art Unit: 2183

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

kv



EDDIE CHAN
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100